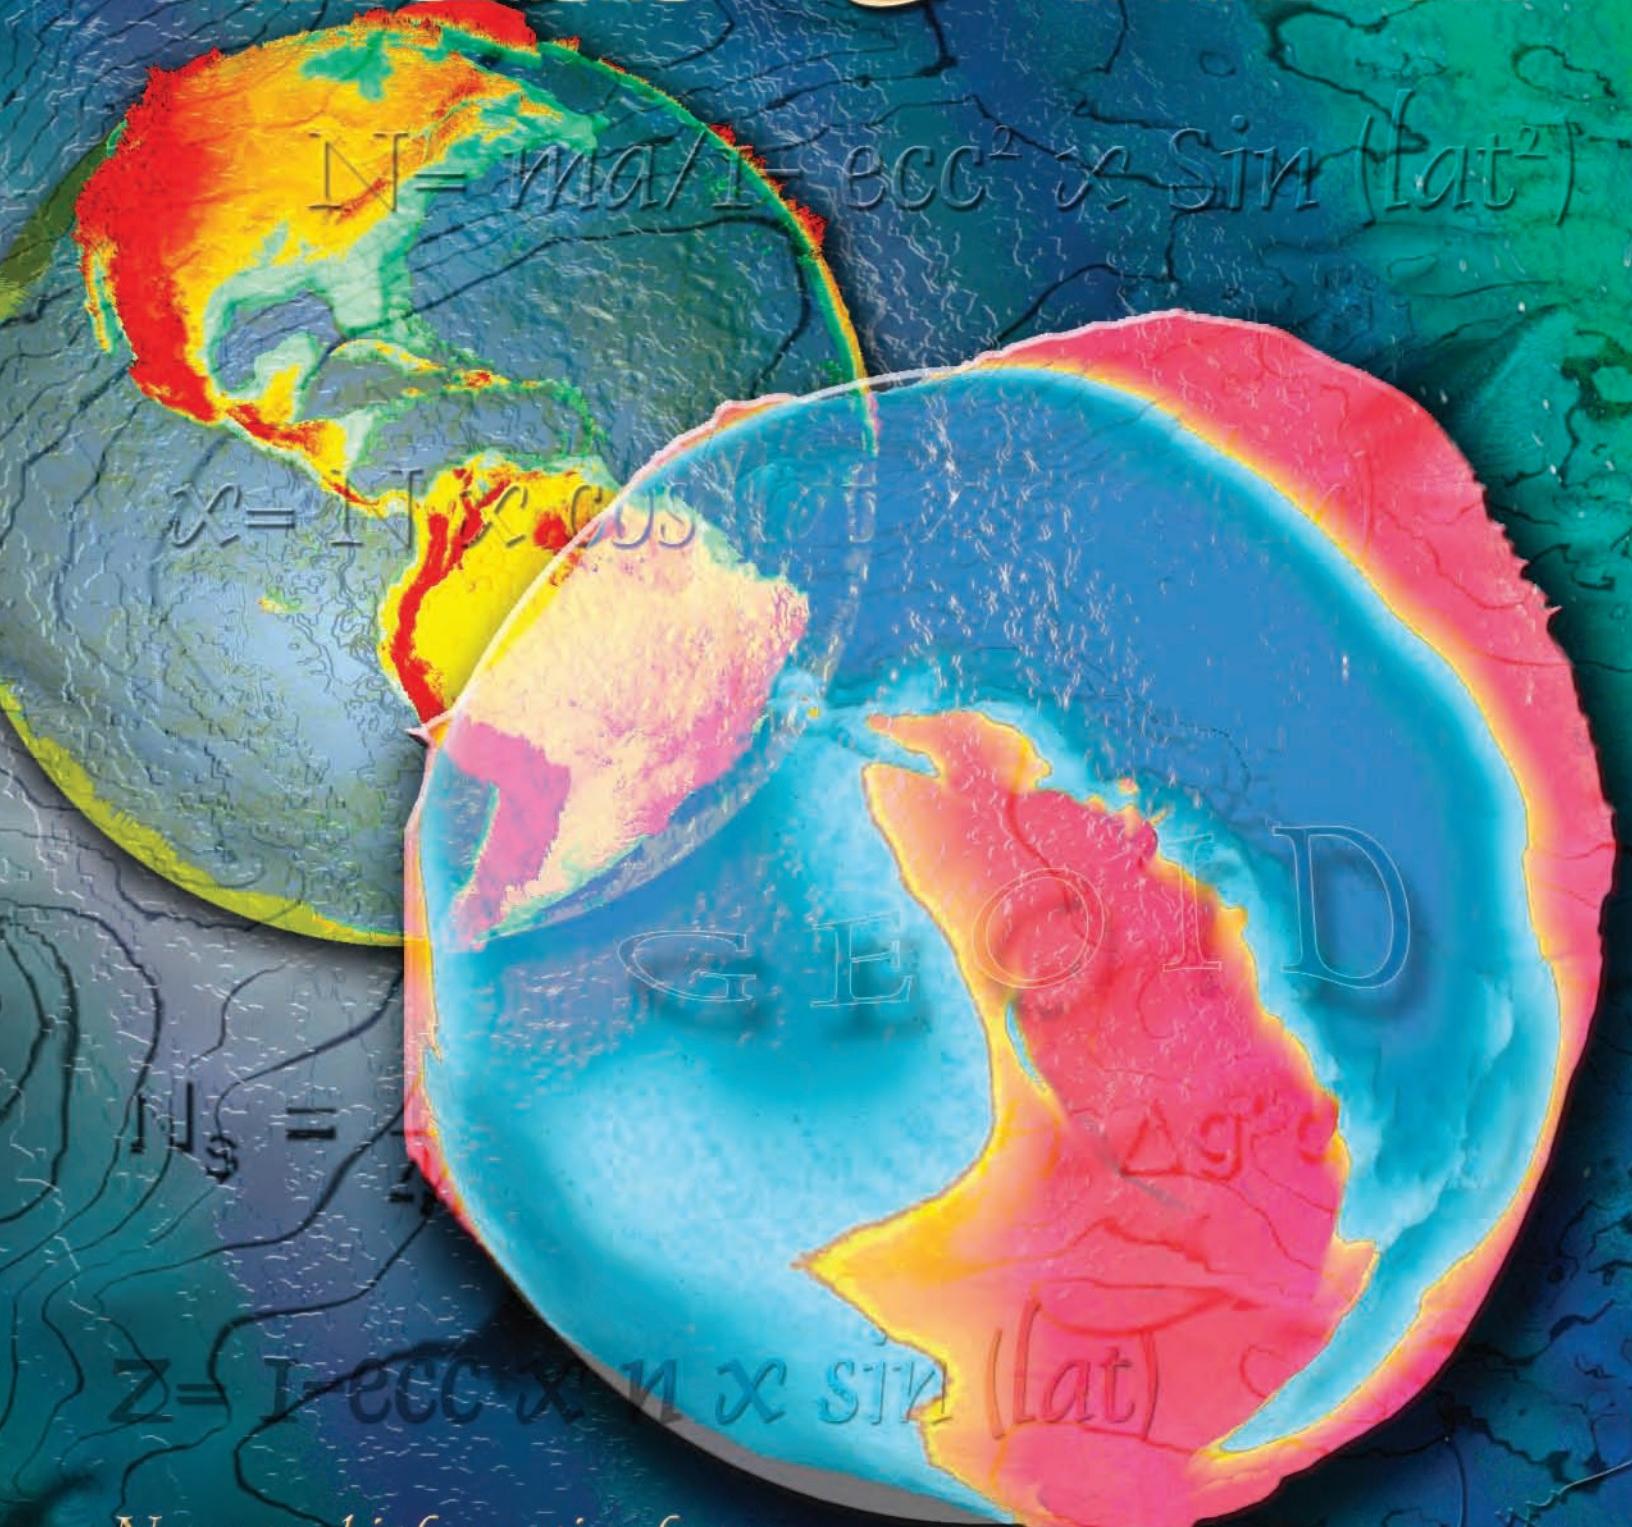




Navigator

NAVO MSRC

SPRING 2006



News and information from...

The Naval Oceanographic Office Major Shared Resource Center

Report Documentation Page			Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.				
1. REPORT DATE 2006	2. REPORT TYPE	3. DATES COVERED 00-00-2006 to 00-00-2006		
4. TITLE AND SUBTITLE NAVO MSRC Navigator. Spring 2006		5a. CONTRACT NUMBER		
		5b. GRANT NUMBER		
		5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S)		5d. PROJECT NUMBER		
		5e. TASK NUMBER		
		5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Oceanographic Office (NAVO),Major Shared Resource Center (MSRC),1002 Balch Boulevard,Stennis Space Center,MS,39522		8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)		
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited				
13. SUPPLEMENTARY NOTES				
14. ABSTRACT				
15. SUBJECT TERMS				
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 28
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified	19a. NAME OF RESPONSIBLE PERSON	

The Director's Corner...



Steve Adamec, NAVO MSRC Director

The NAVO MSRC is undergoing a carefully planned series of enhancements which, when completed by Fall 2006, will provide one of the most capable, productive, and balanced HPC environments ever fielded at this MSRC. These enhancements substantially boost the computational capabilities and resilience of the MSRC across both the classified and unclassified High Performance Computing (HPC) environments we support for the Department of Defense (DoD) HPC Modernization Program (HPCMP).

The most significant enhancements will be the addition of two very large IBM HPC systems, both of which will be based upon IBM's POWER5 processor technology:

- ▶ KRAKEN (the 3000-processor unclassified IBM POWER4+), one of the most successful and requested HPC systems within the NAVO MSRC, will be joined by an unclassified 3000-processor IBM POWER5+ system (named BABBAGE).
- ▶ ROMULUS (the existing 500-processor IBM POWER4+ system) will be transitioned to the classified environment and will be joined by a 1900-processor IBM POWER5+ system (named PASCAL) which will replace the existing classified IBM POWER4 system (MARCELLUS) when it is retired in early Fiscal Year 2007 (FY07).

When all of these upgrades are complete, the effective computing power of the NAVO MSRC will be essentially tripled, as measured by sustainable performance on the HPCMP benchmark suite. All four of these systems will

be configured with two gigabytes of memory per processor, IBM's "Federation" inter-processor switch fabric, and IBM's Global Parallel File System (GPFS), all of which will facilitate the execution of tremendously large applications and also diverse mixes of large applications — applications which typically run as DoD Challenge Projects. To supplement this enormous computational capability, we continue

Enhancing the MSRC to Serve You Better

to enhance and optimize the internal mass storage and networking capabilities of the MSRC for both performance and resilience. Finally, most of you are aware that the HPCMP may be sustaining significant operating budget cuts beginning in FY07. The six HPC centers within the program are slated to receive the bulk of these cuts. Please be assured that our primary goal throughout this budget adjustment process will be the maintenance of a premier HPC environment with the support you have come to expect from all of the centers. We invite you, the DoD user community, to let us continue to assist you in bringing this cutting-edge capability to bear in support of your HPC needs.

Contents

The Naval Oceanographic Office (NAVO) Major Shared Resource Center (MSRC): Delivering Science to the Warfighter

The NAVO MSRC provides Department of Defense (DoD) scientists and engineers with high performance computing (HPC) resources, including leading edge computational systems, large-scale data storage and archiving, scientific visualization resources and training, and expertise in specific computational technology areas (CTAs). These CTAs include Computational Fluid Dynamics (CFD), Climate/Weather/Ocean Modeling and Simulation (CWO), Environmental Quality Modeling and Simulation (EQM), Computational Electromagnetic and Acoustics (CEA), and Signal/Image Processing (SIP).

NAVO MSRC
Code N7
1002 Balch Boulevard
Stennis Space Center, MS 39522
1-800-993-7677 or
msrchelp@navo.hpc.mil

NAVO MSRC Navigator

www.navo.hpc.mil/Navigator

NAVO MSRC Navigator is a biannual technical publication designed to inform users of the news, events, people, accomplishments, and activities of the Center. For a free subscription or to make address changes, contact NAVO MSRC at the above address.

EDITOR:

Gioia Furness Petro, petrogio@navo.hpc.mil

DESIGNERS:

Kerry Townson, ktownson@navo.hpc.mil

Lynn Yott, lynn@navo.hpc.mil

Any opinions, conclusions, or recommendations in this publication are those of the author(s) and do not necessarily reflect those of the Navy or NAVO MSRC. All brand names and product names are trademarks or registered trademarks of their respective holders. These names are for information purposes only and do not imply endorsement by the Navy or NAVO MSRC.

Approved for Public Release
Distribution Unlimited

The Director's Corner

- 2 Enhancing the MSRC to Serve You Better

Feature Articles

- 4 Five Dogs, a Hurricane, and Two IBM POWER5+: A Personal and Professional Katrina Experience
- 5 Drag Reduction by Microbubbles in a Spatially-Developing Turbulent Boundary Layer: Reynolds Number Effect (HPCMP/CAP)
- 9 Practical Parallel Graphics and Remote Visualization
- 12 Mapping the Seabed on an Absolute Reference Frame System Using the Real-Time GIPSY (RTG) Global Differential GPS and RTK Positioning

The Porthole

- 22 Visitors to the Naval Oceanographic Office Major Shared Resource Center
- 24 Boy Scout Jamboree

Navigator Tools and Tips

- 25 Hints for Choosing “Consumable CPUs” and “Consumable Memory” Resource Requests on KRAKEN and ROMULUS

Upcoming Events

- 27 Coming Events

Five Dogs, a Hurricane, and Two IBM POWER5+s: A Personal and Professional Katrina Experience

Dave Cole, Government NAVO MSRC User Support Lead

Articles in the Fall 2005 edition of the Navigator reported the unprecedented effects of Hurricane Katrina on the Naval Oceanographic Office Major Shared Resource Center (NAVO MSRC) and employees. The log that follows, which records my evacuation from the Gulf Coast and eventual return, provides a more personal snapshot of events that is in a small way representative of the ordeals that they experienced and overcame.

But first, a little background. In January 2005 I became the NAVO MSRC Technical Insertion for Fiscal Year 2006 (TI-06) Usability Team Chair. As such, I was charged with keeping the NAVO MSRC Usability Team process on schedule as part of the Center's acquisition of two large IBM POWER5+ clusters through the Department of Defense (DoD) High Performance Computing Modernization Program (HPCMP) TI-06 acquisition effort.

As the NAVO MSRC TI-06 Usability Team Chair, my participation began with a meeting in January 2005 with the TI-06 Performance Team and essentially culminated with a presentation to the TI-06 Collective Acquisition Team (CAT) in November 2005. Keeping the Usability Team process on schedule was especially challenging as Hurricane Katrina struck during the last week of the Phase I analysis.

- ▶ Fri., 26 Aug.: Held first Usability evaluation conference call as scheduled. Departed work under normal hurricane condition status—no preparation needed.
- ▶ Sat., 27 Aug.: Katrina now a threat—spent the day cutting plywood shutters for my house and preparing for evacuation. Liberated two beagles from local kennel for a friend who was out of the

country. This brings our pet count to five small, frisky dogs!

- ▶ Sun., 28 Aug.: Katrina now a Category (CAT) 5 headed for New Orleans/West Mississippi (MS) Coast. Evacuated to Minden, LA (where my parents live). Boarded the frisky dogs.
- ▶ Mon., 29 Aug.: Held second evaluation conference call at 1330 Central Time as scheduled. Welcome surprise that my dad, who avoids technology, has a speaker phone in the kitchen! Used the kitchen table as my command post. Northeast quadrant of Katrina plows into the west MS Coast with a storm surge greater than 20 feet (my house is at a 12-foot elevation). Bought a Universal Serial Bus (USB) compatible keyboard at Wal-Mart for my laptop and began development of the Usability Brief.
- ▶ Tue., 30 Aug.: Moved into Best Western for Digital Subscriber Line (DSL) access. Completed the draft Usability Presentation and passed the ball to Tom Crimmins of the Army Research Laboratory (ARL) at 22:24 so I could begin planning the trek home to pick

Continued Page 21



Five (usually) frisky dogs.

Drag Reduction by Microbubbles in a Spatially-Developing Turbulent Boundary Layer: Reynolds Number Effect (HPCMP/CAP)

Antonino Ferrante and Said Elghobashi, Department of Mechanical and Aerospace Engineering, University of California, Irvine

INTRODUCTION

Experimental evidence during the past three decades indicates that the injection of gaseous microbubbles (diameter ranging from 1 to 1000 microns, and at a relatively large volume fraction (up to $\Phi_v = 0.7$)) into a liquid turbulent boundary layer over a flat plate^{1, 2} or over axisymmetrical bodies³ can reduce the skin friction by as much as 80 percent from its value without bubble injection. However, the basic physical mechanisms responsible for that reduction were not yet fully understood.

This article discusses the physical mechanisms responsible for the reduction of skin friction in a microbubble-laden, Spatially-Developing Turbulent Boundary Layer (SDTBL) over a flat plate⁴, and the effects of increasing Reynolds number on drag reduction.⁵ This discussion is based on the results of the author's Direct Numerical Simulations (DNS) of a microbubble-laden SDTBL. These simulations were performed on High Performance Computing (HPC) highly scalable supercomputers (CRAY T3E and IBM Power4+ (KRAKEN)) at the Naval Oceanographic Office (NAVOCEANO) Major Shared Resource Center (MSRC).

MATHEMATICAL DESCRIPTION

Figure 1 shows a schematic of the SDTBL flow where the gravitational acceleration vector is perpendicular to the wall and pointing downward. The DNS used in this simulation employs the Eulerian-Lagrangian approach to solve the fluid continuity and momentum equations in an Eulerian

framework. The bubble acceleration equation, on the other hand, is solved for each bubble to track its trajectory in time.^{4, 5}

The governing equations of the fluid motion account for the instantaneous local volume fraction of the bubbles. The bubble equation of motion includes terms representing the added mass, carrier fluid inertia, Stokes drag, buoyancy, and lift force. The governing equations^{4, 5} were discretized in space using a second-order finite difference scheme—except for the mean advection terms, which were evaluated via a fifth-order upwind differencing scheme.

Time integration in this simulation was performed using the second-order Adams-Bashforth scheme. The discretized Poisson equation for pressure was solved using a cosine transform in the streamwise direction, a Fast Fourier Transform (FFT) in the spanwise direction, and Gauss elimination in the wall-normal direction. The discrete cosine and Fourier transforms were computed using the Fastest Fourier Transform in the West (FFTW) C subroutine library.⁶

Continued Next Page...

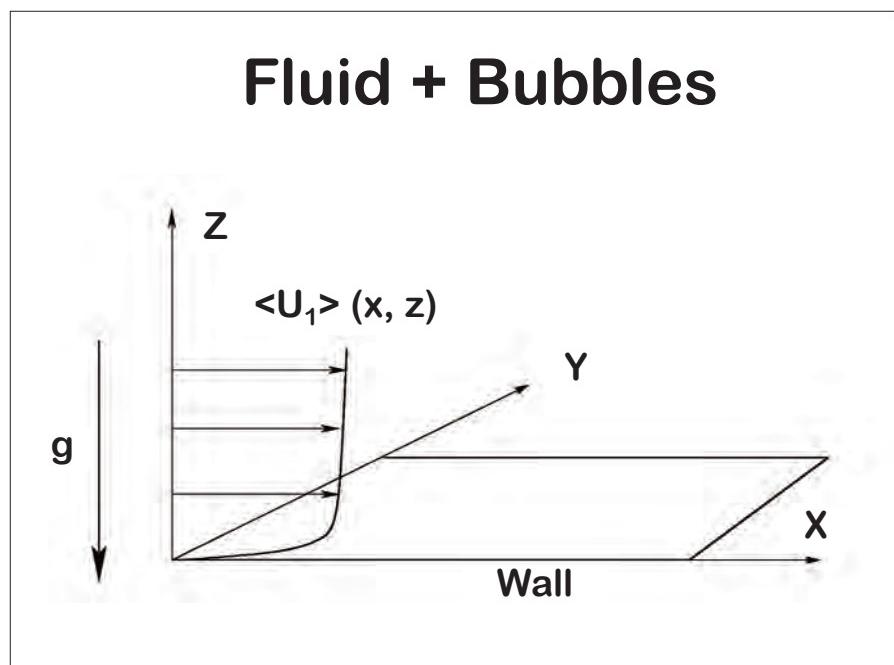


Figure 1. Schematic of microbubble-laden turbulent boundary layer flow over a flat wall.

CAP PARALLEL PERFORMANCE TESTS

During the past four years the authors have used their newly-developed parallel code (DNSBLB, written in FORTRAN 90/MPI), which performs a DNS of a microbubble-laden SDTBL.^{4, 5, 7, 8, 9} DNSBLB is parallelized with a One-Dimensional (1D) domain decomposition in the spanwise y -direction (j) of the three-dimensional computational domain.* DNSBLB was originally written for the CRAY T3E. During the Capability Applications Project (CAP) 2004, the CRAY T3E version of DNSBLB was converted to run on KRAKEN. The IBM version of

DNSBLB is more than ten times faster than its T3E version. The scalability runs of the DNS code were performed using up to 1024 processors for two different computational meshes: a coarse mesh of $256 \times 256 \times 96 = 6 \times 10^6$ grid points, and a fine mesh of $1024 \times 1024 \times 96 = 101 \times 10^6$ grid points; and for four different numbers of bubbles 100 thousand, 16 million, 100 million, and 200 million.

The scalability tests were performed in a one-way coupling regime, i.e., no effects of bubble on turbulence were accounted for. The two-way coupling increases the Central Processing Unit (CPU) time but does not worsen the

scalability of the code. Figure 2 shows the CPU time in seconds per processor, per time step needed for the integration of the governing equations of both the fluid and bubbles, versus the number of processors used on KRAKEN.

Different lines correspond to different computational meshes and numbers of bubbles. The CPU time monotonically decreases as the number of processors increase (see Figure 2) for all the tests performed.

Figure 2 also shows that for the coarse mesh with 100 thousand bubbles, the slope of the line (CPU time versus number of processors) is close to -0.1, whereas for the fine mesh with 200 million bubbles, the slope is close to -1.

The slopes for the other tests are between -0.1 and -1. This means that the performance of the DNS code improves as the number of grid points of the computational mesh and the number of bubbles simulated increase. Furthermore, for the fine mesh the code shows a nearly ideal scalability since the line in Figure 2 has a slope close to -1.

DRAG REDUCTION BY MICROBUBBLES: REYNOLDS NUMBER EFFECT

The DNS results⁴ for the microbubble-laden SDTBL for $R_e \Theta = 1430$ with volume fraction ranging from $\Phi_v = 0.001$ to 0.02 show that the presence of bubbles results in a local positive divergence of the fluid velocity, $\nabla \cdot U > 0$. This creates a positive mean velocity, $\langle U_3 \rangle$, normal to (and away from) the wall which, in turn, reduces the mean streamwise velocity and displaces the quasi-streamwise longitudinal vortical structures away from the wall as in Figure 3. This displacement has two main effects:

- ▶ It increases the spanwise gaps between the wall streaks associated with the sweep events and reduces the streamwise

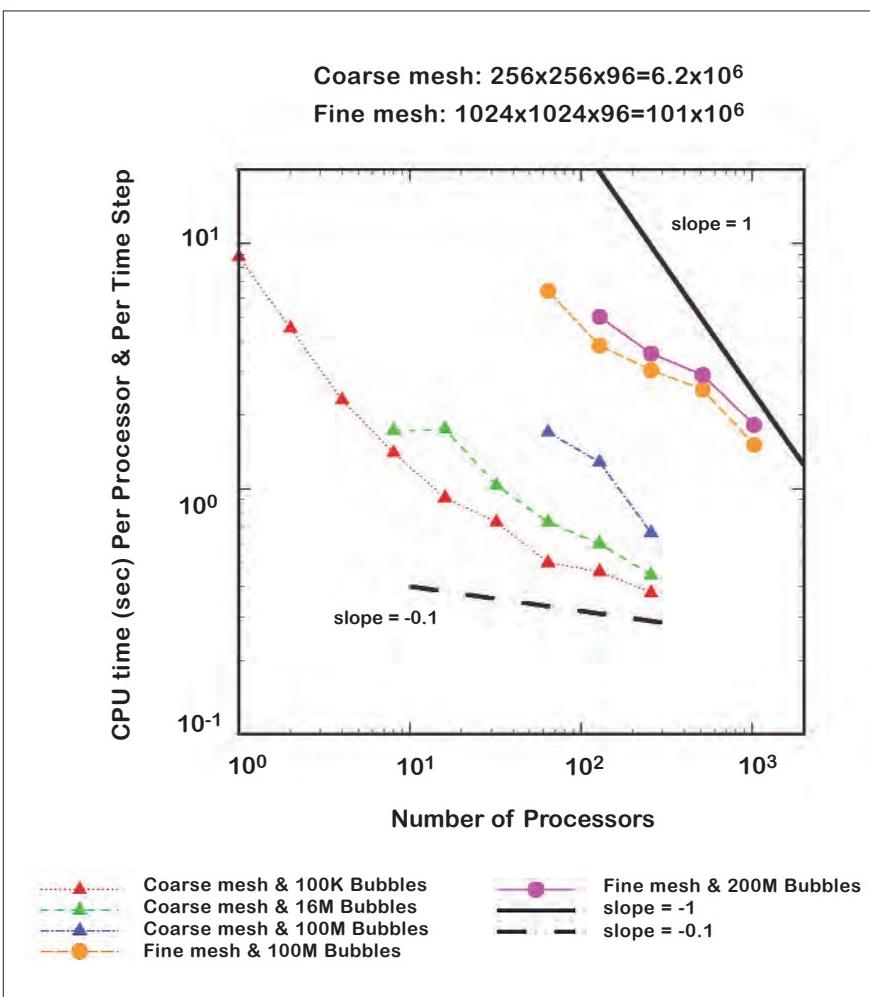


Figure 2. Scalability on IBM P4+ (KRAKEN) of DNS code.

*If np is the number of processors and N , even multiple of np , is the number of grid points in the j direction, then each processor makes calculation on $N=np$ planes ($x-z$ planes) of the computational domain.

velocity in these streaks, thus reducing the skin friction.

- It moves the location of peak Reynolds stress production away from the wall to a zone of a smaller transverse gradient of the mean streamwise velocity (i.e., smaller mean shear), thus reducing the production rate of turbulence kinetic energy and enstrophy.

During CAP 2004 the authors were able to simulate the microbubble-laden SDTBL for $R_e\Theta = 2900$ with volume fraction $\Phi_v = 0.01$ and bubble diameter of $40\mu m$ using a computational mesh of $1024 \times 512 \times 128 = 67 \times 10^6$ grid points and 29 million bubbles.

The computations were performed on 512 processors of the KRAKEN system. The bubble-laden flow simulations required 9.6 CPU hours

per processor and 124 gigabytes of maximum total memory.

The DNS results⁵ show that increasing Reynolds numbers from 1430 to 2900 decreases the percentage of drag reduction from 22 to 19 percent. Increasing $R_e\Theta$ “squeezes” the quasi-streamwise vortical structures toward the wall, whereas microbubbles “push them away” from the wall.

The net result of these two opposing effects determines the amount of skin friction reduction by the microbubbles.

The displacement action by the microbubbles is a result of the local positive velocity divergence, $\nabla \cdot U$, created by their concentration gradients. Thus, the volume fraction of bubbles that is responsible for the reduction of skin friction in a low Reynolds number SDTBL is not sufficient to produce

the same amount of reduction in skin friction at higher Reynolds number.

SUMMARY

This article has briefly discussed the physical mechanisms of drag reduction by microbubbles and the Reynolds number effect.^{4, 5}

Furthermore, it reports some details on the simulations of parallel DNS code and its scalability on the NAVOCEANO MSRC KRAKEN system.

In conclusion, the CAP program has considerably helped the author's drag reduction by microbubbles research to explain the effect of Reynolds number on drag reduction by allowing the use of a large number of processors and extended CPU hours which were not allowed in the “standard” queues.

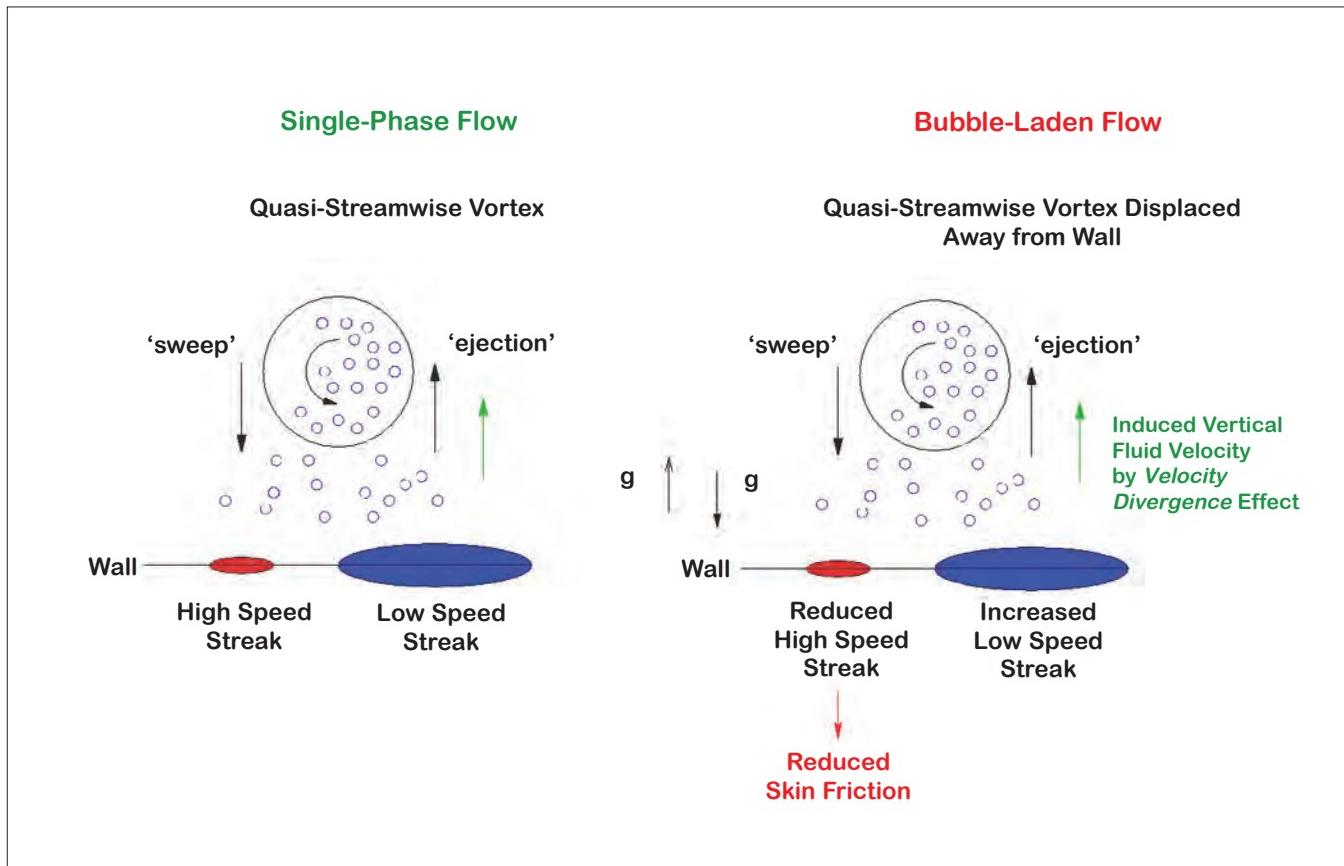


Figure 3. Schematic of the drag reduction mechanism in a microbubble-laden SDTBL.

Acknowledgements

This work was supported by Office of Naval Research (ONR) Grant No. N00014-05-1-0059 under the direction of Dr. Patrick Purtell. The computations were performed on the IBM Power4+ (KRAKEN) located at the NAVO MSRC at the John C. Stennis Space Center (Mississippi), and on CRAY T3E (MD) located at the U.S. Army High Performing Computing Research Center (AHPCRC, Minnesota). The authors thank the outreach group at the NAVOCEANO MSRC and the support group at AHPCRC for allowing jobs with CPU time limits larger than those allowed in the “standard” queues.

References

1. Madavan, N.K., S. Deutsch, and C.L. Merkle, “Reduction of Turbulent Skin Friction By Microbubbles,” Physics of Fluids, 27: 356-363, 1984.
2. Pal, S., C. L. Merkle, and S. Deutsch, “Bubble Characteristics and Trajectories in a Microbubble Boundary Layer,” Physics of Fluids, 31(4): 744-751, 1988.
3. Deutsch, S., and S. Pal, “Local Shear Stress Measurements on an Axisymmetrical Body in a Microbubble-Modified Flow Field,” Physics of Fluids, 2: 2140-2146, 1990.
4. Ferrante, A., and S. Elghobashi, “On the Physical Mechanisms of Drag Reduction in a Spatially Developing Turbulent Boundary Layer Laden with Microbubbles,” Journal of Fluid Mechanics, 503: 345-355, 2004.
5. Ferrante, A., and S. Elghobashi, “Reynolds Number Effect on Drag Reduction in a Microbubble-Laden Spatially Developing Turbulent Boundary Layer,” Journal of Fluid Mechanics, 543: 93-106, 2005.
6. Frigo, M., and S.G. Johnson, “The Design and Implementation of FFTW3,” Proceedings of the IEEE, 93(2): 216-231, 2005.
7. Ferrante, A., and S. Elghobashi, “A Robust Method for Generating Inflow Conditions for Direct Simulations of Spatially-Developing Turbulent Boundary Layers,” Journal of Computational Physics, 198:372-387, 2004.
8. Ferrante, A., S. Elghobashi, P. Adams, M. Valenciano, and D. Longmire, “Evolution of Quasi-Streamwise Vortex Tubes and Wall streaks in a Bubble-Laden Turbulent Boundary Layer Over A Flat Plate,” Physics of Fluids, 16(9): S2, 2004.
9. A. Ferrante, “Reduction of Skin-Friction in a Microbubble-Laden Spatially Developing Turbulent Boundary Layer Over a Flat Plate,” PhD thesis, University of California, Irvine, 2004.

Practical Parallel Graphics and Remote Visualization

Sean Ziegeler, Visualization Software Engineer, NAVO MSRC VADIC

The purpose of scientific visualization is, quite often, to analyze the output of a computational model. As computational power increases, model output grows in size, which introduces new obstacles that prevent users from analyzing their data—especially if they are not collocated with the computing resources.

The traditional approach to scientific visualization from remote sites has been to transfer the output data from the computational system to the user's local site as shown in Figure 1. This is less feasible with larger data because it may be too large to (1) transfer over the network between sites in a reasonable amount of time, (2) fit on the user's system disk space, or (3) render on the user's local graphics computer.

To overcome all three problems, a combination of techniques must be utilized: computational data access, parallel graphics, and remote visualization.

Computational data access involves the transfer of the data between a computational and a visualization system, either by staging it on large

disks before visualization or transferring it “on-the-fly” during visualization using a remote file system.

The next technique, parallel graphics, uses large graphics systems (such as clusters) to render in parallel fashion by distributing the load across multiple Graphics Processing Units (GPU's).

Finally, remote visualization can deliver the rendered images, reduced now merely to pixels (and much smaller than the original data), from the graphics system to the user's local system. This new approach is shown in Figure 2.

REMOTE DATA ACCESS

Staging the data on a visualization system is very similar to transferring the data from the computational system to the user's local system. This is the preferred method if the data set is to be visualized repeatedly, is not too large to fit on the visualization system, and requires high performance access.

The result of staging the data in parallel on a graphics cluster is

improved performance. This improvement is because a parallel file system consists of multiple disks spread out over many or all nodes of the cluster. An example at the Naval Oceanographic Office (NAVOCEANO) Major Shared Resource Center (MSRC) is the /scr file system on the visualization cluster, SEAHORSE.

If the data will not fit on /scr, will not be visualized repeatedly, or won't require the high-speed access of a parallel file system, then the Secure Shell File System (SSHFS) is an excellent tool for accessing data on a separate system as if it were local.

SSHFS uses the Secure Shell (SSH) program that is already installed on all MSRC systems to facilitate authenticated and encrypted file access.

The difference between SSHFS and a traditional file transfer is that the remote files will appear to be like any other file on the local system. Command Example 1 shows how to “mount” or begin, briefly utilize, and end an SSHFS session.

Continued Next Page...

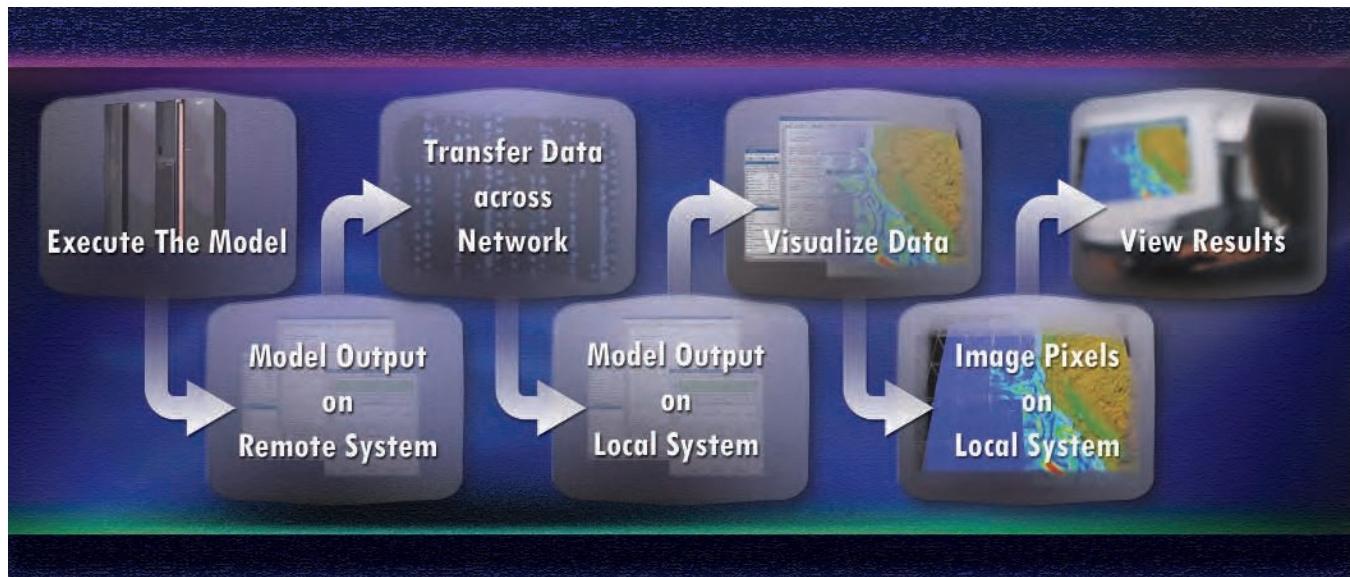


Figure 1. Traditional scientific visualization by transferring data to user's site.

COMMAND	PURPOSE
jules> ls /u/home/myname	List the files in home directory on JULES.
file1 file2 file3	(Results)
seahorse> mkdir /u/home/myname/myjules	Create a “mount point” on SEAHORSE.
seahorse> sshfs jules:/u/home/myname /u/home/myname/myjules	Connect to JULES.
seahorse> ls /u/home/myname/myjules	Files should match those on JULES.
file1 file2 file3	(Results)
seahorse> fusermount -u /u/home/myname/myjules	Disconnect from JULES when finished.

Command Example 1. How to mount, briefly utilize, and end an SSHFS session.

Performance is usually better if multiple mounts from multiple nodes are used when executing a parallel visualization application.

For security reasons, this requires multiple SSHFS mounts from the master nodes to the destination system and then from the compute nodes to the master nodes. Figure 3 illustrates the initial decrease in bandwidth due to the overhead of SSHFS and the increase in bandwidth achieved overall. Each type of corresponding connection is conceptually diagrammed as well. To make this complex process easier, the NAVOCEANO MSRC Visual Analysis and Data Interpretation Center (VADIC) has implemented a script that automates the creation of the mount points. The following commands use that script to establish a connection from two master nodes (1 and 2) and four compute nodes (3 through 6) of SEAHORSE to two corresponding master nodes of KRAKEN (1 and 2):

```
seahorse> psshfs -mount
kraken1:/u/home/myname
/u/home/myname/mykraken 1 3
4 -mount kraken2:/u/home/
myname /u/home/myname/
mykraken 2 5 6
```

PARALLEL GRAPHICS

For basic visualization needs and the do-it-yourself user, Paraview is a useful, general-purpose visualization application. With a little up front learning, a user can start from scratch or have VADIC support staff create an initial visualization that can be extended later by the user.

Paraview's true power lies in its parallelism. It can be used in conjunction with the Message Passing Interface (MPI) in a way that is mostly transparent to the user but takes advantage of the multiple nodes of a graphics cluster. The following command runs the MPI version of Paraview on SEAHORSE using four processors on nodes 2 through 5:

```
seahorse> mpirun_rsh -np 4
scn2 scn3 scn4 scn5
DISPLAY=:0.0 /usr/local/
paraview-2.4.1-mpi/bin/
paraview --RenderModule=
MPIRenderModule
```

The previous command runs Paraview locally on SEAHORSE. To use Paraview from an off-site location, see the Remote Visualization section of this article.

Unfortunately, Paraview is not as well-suited to tasks such as analyzing time-

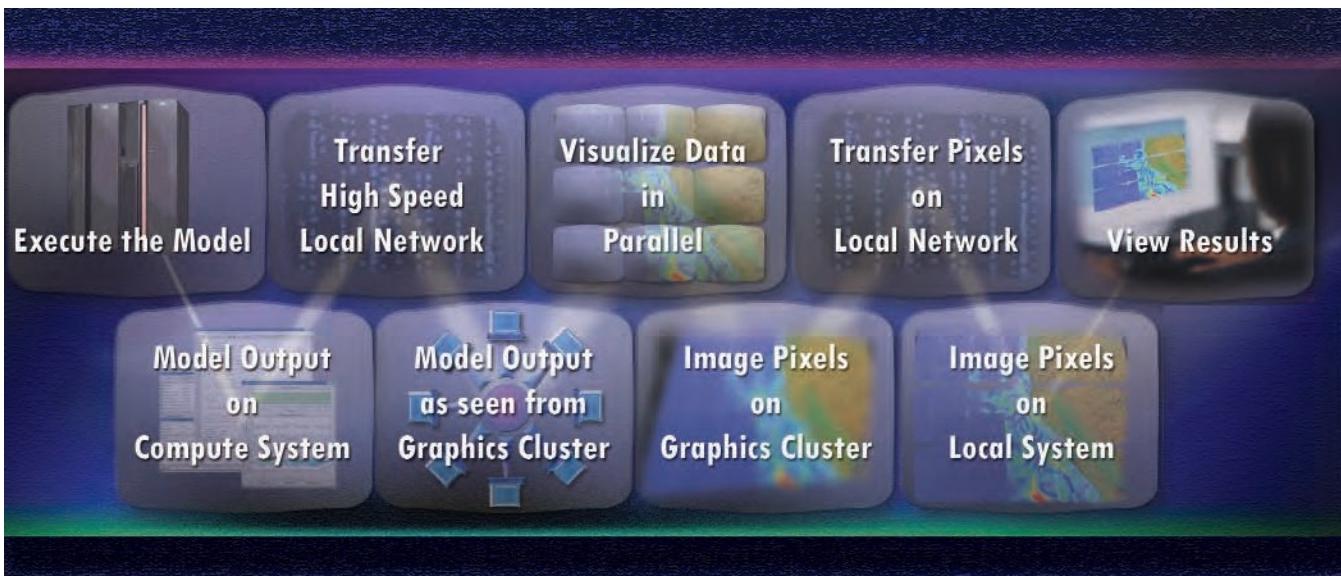


Figure 2. Revised process with computational data access, parallel graphics, and remote visualization.

variant fluid flow and direct volume rendering of structured data sets. Furthermore, it is not as easy to modify as it is to use. While a custom visualization application (which is typically faster and simpler to use) can be developed for specific data to overcome this, the disadvantage is that it requires a visualization expert to design and can potentially take more up-front development time. One way to make development of a custom application easier is to use MPI. Though a combination of graphics and parallel programming may sound like a daunting task, it is not as difficult as might be imagined. The more advanced MPI communication primitives, such as broadcasts and reductions, while not required, allow for even simpler and more elegant parallel rendering. Finally, these advanced communication primitives assist when developing a user interface, for example, by broadcasting user events like mouse clicks to all of the rendering processes. However, alternative approaches to MPI exist for custom applications. The VADIC is currently evaluating two such options: Chromium and Virtual Graphics Platform (VGP). Chromium is an open-source library that distributes

parts of the graphics to different processes, while VGP is a commercial software with a similar function.

REMOTE VISUALIZATION

Remote rendering is not a novel concept. In fact, it has been possible on UNIX systems for many years over a remote X-Windows connection.

What remote X-Windows rendering does is force the local system (versus the remote rendering system) to render all of the graphics commands. While this works well sometimes, it is typically the case that large data sets require a proportionally large amount of rendering which can overwhelm the local system.

An approach that better takes advantage of a graphics cluster is to do all of the rendering on the remote system and merely send the visualization's pixels to the local system. This feature is available in Paraview and, coupled with Paraview's ease of use and parallel capability, makes it yet more attractive as a visualization tool.

To use this feature, Paraview must be executed in client/server mode. Command Example 2 shows a simple example of a client and server process executing on two example systems.

Note that the client system, "computer2" in this example, must have Paraview installed. To run Paraview in parallel the commands are similar to that discussed in the Parallel Graphics section of this article. Note that the commands (which follow) are slightly different than before.

Also, it is important to realize that most users will need to properly configure their accounts to use Command Example 2. Users should see the "More Information" section at the end of this article for assistance. The NAVOCEANO MSRC VADIC has also created a method to remotely display custom visualization

Continued Page 20

COMMAND	PURPOSE
computer1> pvserver	Computer1 will host the server process
computer2> pvclient - sh=computer1	Client on Computer2 is instructed to contact Computer1

Command Example 2. A simple example of a client and server process executing.

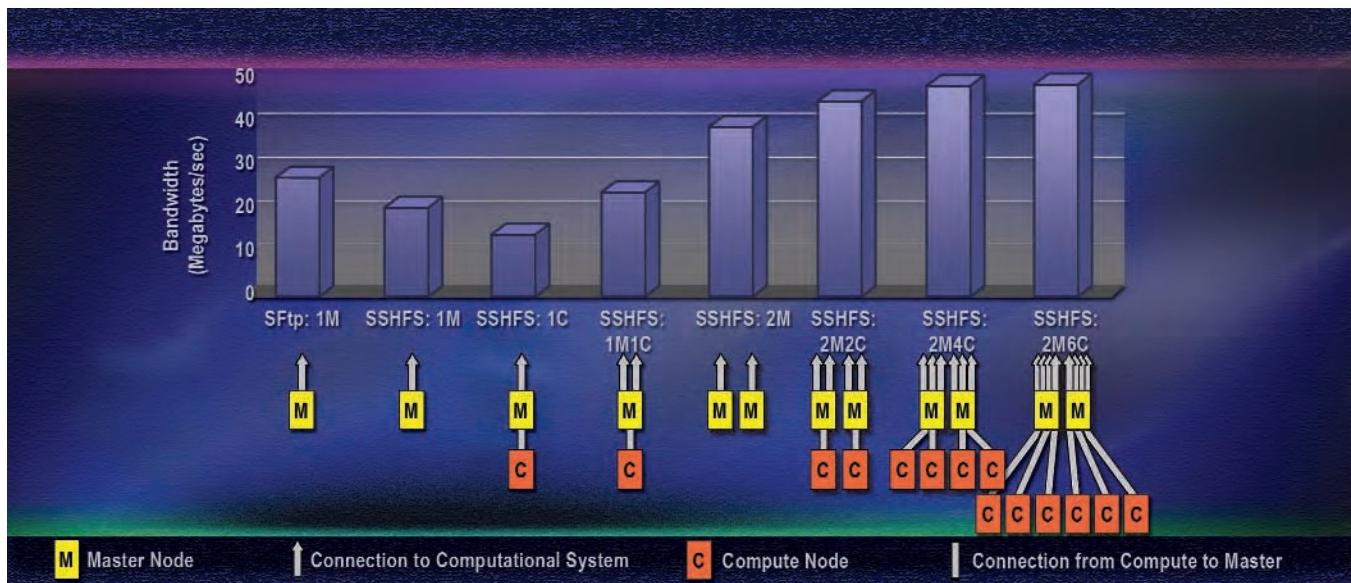


Figure 3. Bandwidth of single SFTP and various single and parallel SSHFS configurations.

Mapping the Seabed on an Absolute Reference Frame System Using the Real-Time GIPSY (RTG) Global Differential GPS & RTK Positioning

Elliot N. Arroyo-Suarez, Naval Oceanographic Office

Maxim F. van Norden, Naval Oceanographic Office

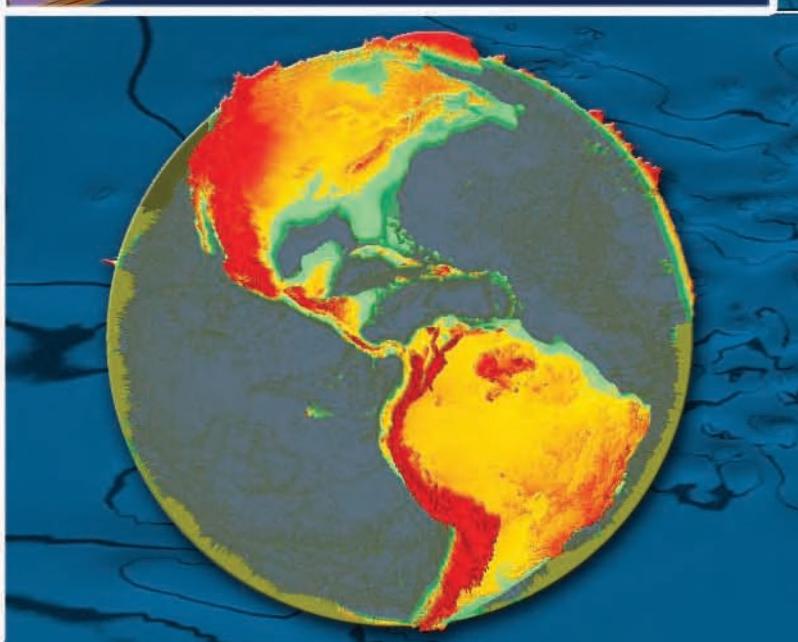
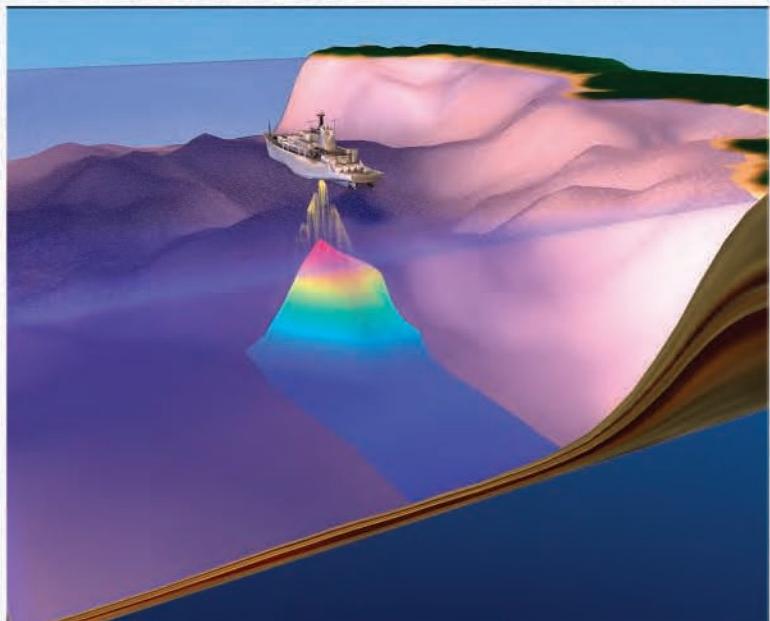
Chad Saxon, Visualization Software Engineer,

NAVO MSRC Visual Analysis and Data Interpretation Center

OVERVIEW

Though the Naval Oceanographic Office (NAVOCEANO) has 175 years of hydrographic and navigation experience, it can no longer rely on the traditional methods for conducting hydrographic surveys. The terrorism threat and nations unwilling to provide access to shore sites for hydrographic survey stations has severely decreased traditional shore-dependent hydrographic survey operations. In addition, the NAVOCEANO focus has shifted to providing the warfighter with high-resolution, near-real-time depiction of the battlespace environment. This shift in focus requires the collection and fusion of current sensor data with historic data to produce the best possible area charts and model/simulation forecasts for battlespace awareness.

To improve navigation and depth measurement positioning accuracy, NAVOCEANO recently acquired Global Positioning System (GPS)



receivers from NavCom Technology, Inc., with Real-Time Kinematic (RTK) and StarFire™ positioning mode capabilities. StarFire™ is a commercialized version of the NASA Jet Propulsion Laboratory (JPL) Real-Time GPS-Inferred Positioning System (GIPSY) (or RTG) and Orbit Analysis Simulation Software (OASIS) Global Differential GPS (GDGPS) corrections. These RTG techniques and corrections are installed aboard NAVOCEANO survey platforms. The NavCom GPS systems are also employed by the Compact Hydrographic Airborne Rapid Total Survey (CHARTS), Fleet Survey Team (FST), and International Surveys Program (HYCOOP) survey teams with their unique hydrographic survey systems.

The RTG derived corrections (also referred as StarFire™) produce globally uniform precise GPS orbit and clock corrections. Using the StarFire satellite-based RTG correctors with corrections for Ionospheric and Tropospheric delays computed at the receiver site and solid earth

tide, NAVOCEANO tests indicate that International Hydrographic Organization (IHO)¹ horizontal and vertical accuracy standards can be achieved without establishing shore stations.

These accuracies in the horizontal and vertical, along with techniques to measure water levels with RTG/Real-Time Kinematic (RTK) Positioning and Telemetry Buoys (P&TB), indicate a capability to map the seabed on a seamless geocentric reference system anywhere in the world to IHO Order I standards.

Where higher accuracy standards are required, the system is also capable of measuring the seabed to IHO Special Order accuracy standards with the use of RTK and RTK Extend near RTK shore-based stations.

This article discusses the development of a new Concept of Operations (CONOPS) at NAVOCEANO using technology to accurately position satellites in real-time, now applied to support the Navy's need for high-accuracy navigational charting and battlespace characterization anywhere in the littorals (outside territorial waters). Furthermore, the intention of this CONOPS is to employ NAVOCEANO assets to impact operational and tactical time scales.

U.S. NAVY DOCTRINE AND MILITARY SURVEYS

The environmental characterization of the battlespace is called Intelligence Preparation of the Environment (IPE). The aim of this characterization is to add value and knowledge to the information used by an operational commander and do it within that military operations' decision-making cycle.

With these new challenges, traditional hydrographic surveys can no longer reliably provide timely or accurate data for IPE.

NAVOCEANO hydrographic data are usually acquired with the cooperation

or active participation of the host nation; however, for many areas of the world recent hydrographic data do not exist or may not be available to U.S. planners for political or military reasons.

In many of these cases it will not be possible to establish shore stations or the security risk is too high to justify deploying survey teams ashore for geodetic or tidal data collection.

Rapid Environmental Assessments (REA) that include hydrographic operations are required to describe and understand the battlespace environment. This will require the collection and fusion of on-scene, remotely sensed and historic data to produce the best possible description of the operating environment.

Figure 1 approximately depicts the Exclusive Economic Zone (EEZ) areas of the world and their possible expanded areas. The United Nations Convention on Law of the Sea (UNCLOS) gives the Coastal State jurisdiction over Marine Scientific Research (MSR) in the EEZ.

The U.S. Government view, however, is that UNCLOS MSR provisions do not apply to military survey activities in the EEZ.² With the acquisition of

StarFire™ (and its ability to map the seabed at a decimeter level without shore installations) any conflict created by these differing interpretations of UNCLOS MSR jurisdictions is avoided.

Traditionally, NAVOCEANO employed the following methods to try to achieve IHO horizontal accuracy requirements:

- ▶ U.S. Coast Guard or host nation Differential GPS (DGPS) beacons.
- ▶ WADGPS - A worldwide subscription service that employs a network of ground stations whose pseudo-range correctors are relayed through the International Marine Satellite (INMARSAT) system. NAVOCEANO personnel have observed horizontal errors of up to 14 meters (95 percent confidence) produced from the current commercial WADGAPS correction services.
- ▶ NAVOCEANO-established DGPS pseudo-range corrector station at a geodetically established point.

In order to correct each depth measurement to a vertical datum to IHO standards, water level or tides corrections need to be generated. The process of generating such depth correctors involves the analysis of tide phases and tide amplitudes above short-term mean sea level. These

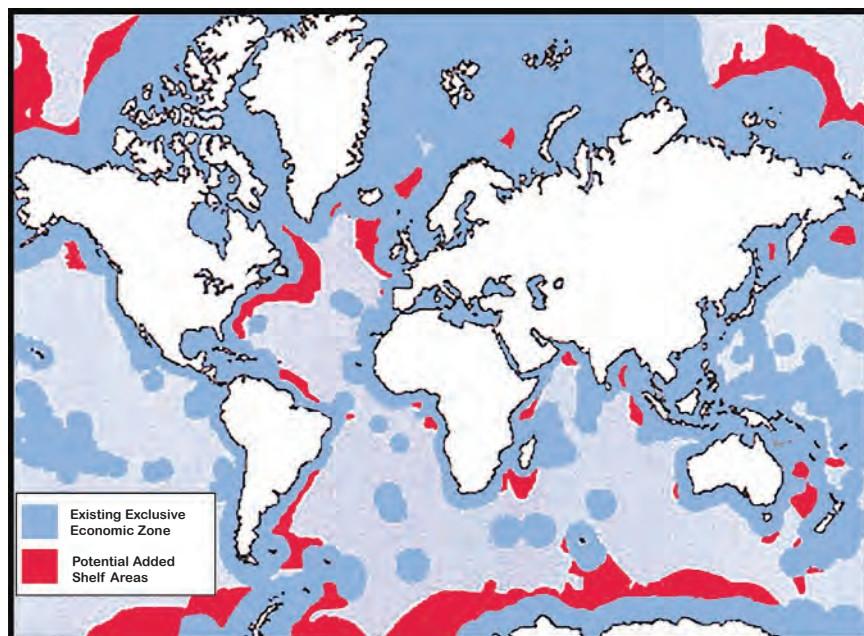


Figure 1. Approximate EEZ areas of the world.

water levels are determined by constructing co-tidal or co-range charts or by using numerical models such as ADCIRC, FES-99, and CU-Tides.

Co-tidal charts are tide constituent models that are determined by inferring the major tide constituents (i.e., phase lag and amplitude) from neighboring established tide stations. Often, areas with unknown constituents may be filled in with virtual tide stations using constituents from numerical models. Each constituent is then gridded for phase lag and amplitude throughout the survey area.

The Z_{00} value is then determined at each station by datum transfer from a known primary tide station, being careful to transfer the datum only between stations with the same tide characteristic. (See Equation 1.) Finally, the area is divided into tide zones, with each zone having a set tolerance (usually +/- 10 to +/- 20 centimeters) from the average set of constituents and Z_{00} value for each zone.

The hydrographer needs to know the MSL at the local gauge and the tide range for this equation. As little as 50

hours of data for semidiurnal tides, if taken at Spring tide, can give sufficient results. Diurnal tides require 15 days of observation to resolve the four primary constituents (i.e., m, r/R, Z_{00} , and S.C.).

If "m" is the height of observed MSL above the WGS-84 ellipsoid (i.e., if gauge zero is the WGS-84 ellipsoid), then "d" is the distance Chart Datum is above the WGS-84 ellipsoid at the secondary station. Figure 2 depicts a typical hydrographic survey scenario with a cross section of tidal zones and water level correctors to reduce the depth to a chart datum.

In order to achieve IHO standards for vertical accuracy, the NAVOCEANO, the HYCOOP program, and the FST conduct RTG and RTK GPS hydrographic surveys using a combination of RTK base station (established at a nearby coastal location) data and, in other cases, just RTG positioning without the need for land-based stations.

The distance from the WGS-84 reference ellipsoid to the chart datum

Chart Datum (d) above gauge zero at local (secondary) station is given by:

$$d = m - (S.C.) - (r/R) Z_{00}.$$

Where:

- ◆ m = height of observed Mean Sea Level (MSL) above gauge zero at local station
- ◆ r/R = ratio of local station/primary station tide ranges for semi-diurnal type tides
- ◆ r/R = ratio of the sum of the M2+S2+K1+01 constituents at each station (local station/primary station) for diurnal type tides [EAS]
- ◆ Z_{00} = height of long term MSL above chart datum at the primary station
- ◆ S.C. = Seasonal Change correction is the difference between long term MSL and observed MSL at the primary tide station.

Continued Next Page...

Equation 1

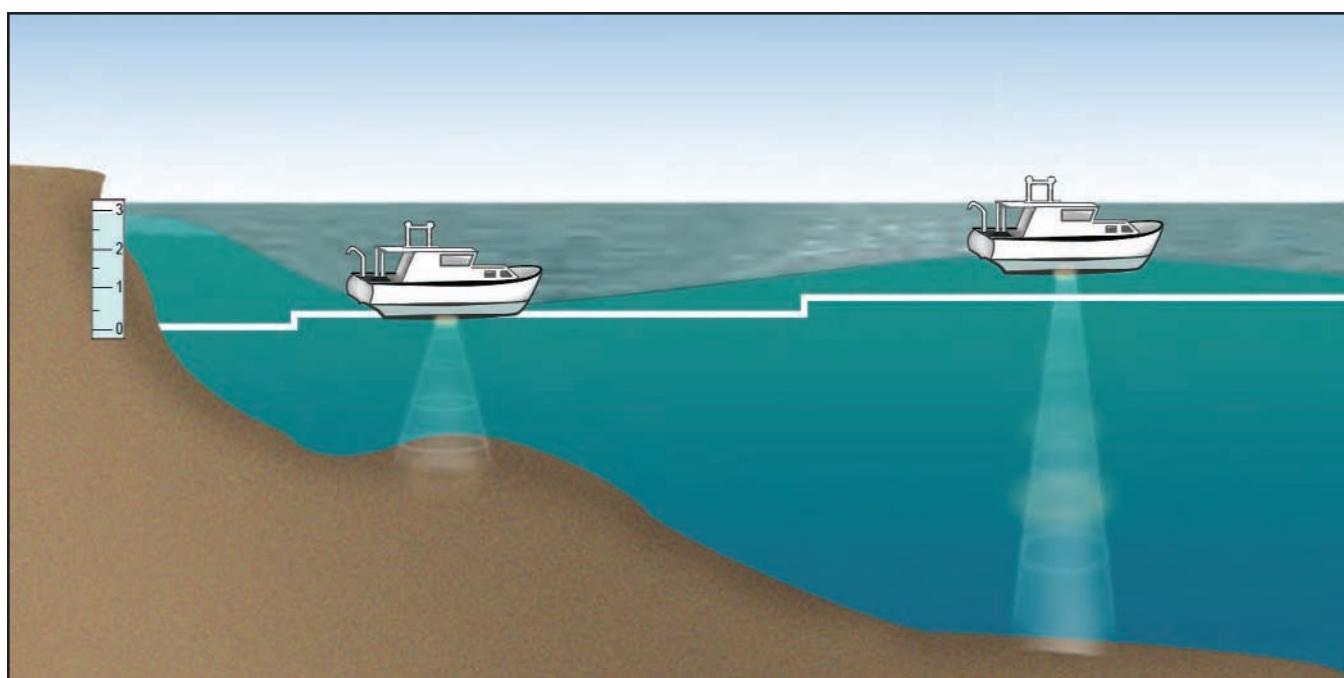


Figure 2. Cross section of Tidal zoning analysis (tidal phase and amplitude shifts).

is determined by placing a tide gauge alongside a GPS base station. This results in a chart datum referenced to the seamless Earth Centered Earth Fixed (ECEF) geocentric WGS-84 reference system.

The following expressions (and Figure 3) explain how Ellipsoidal Depths (ED) and Chart Soundings (CS) products are determined:

- ▶ ED = h_Nav + Measured Depth (D)*
- ▶ CS = ED - SEP
- ▶ SEP = Separation Between the Chart Datum and the WGS-84 Ellipsoid

PERFORMANCE TESTING OF THE STARFIRE™ -BASED GPS SYSTEM

In 2002, NAVOCEANO began evaluating state-of-the-art GPS technologies and, in particular, the implementation of technologies and

techniques to derive water level corrections to bathymetric soundings utilizing RTK and RTG GPS systems.

The RTG method produces globally-uniform precise GPS orbit and clock corrections. Ionospheric and multipath corrected range measurements are performed at the receiver site by processing the L1 and L2 code and carrier measurements.

In addition, a constrained estimate of the tropospheric refraction is made using the data from all the GPS Satellites.³ The GPS system could be utilized in RTK mode whenever the mission requires centimeter-level accuracy and secure access to land can be obtained.

The capability of storing GPS raw observables was also an essential requirement. Since the RTK solutions are limited in range, a Post-Processed Kinematic (PPK) position would be computed.

SUMMARY OF EVALUATION RESULTS

STATIC TESTS

Figure 4 shows in orange an uncorrected RTG position solution without the application of the time-dependent fluctuation of earth crust flex due to the earth tide. Note the Solid Earth Tide (SET) correction in Stennis Space Center (SSC) is approximately 50 centimeters in amplitude over that period of time.

The blue plot shows an RTG position solution with the IERS Note 21 SET model corrector applied, rendering a vertical position solution of approximately 24 centimeters at 95 percent Confidence Level. This finding was so encouraging that it led to this new paradigm of conducting military hydrographic surveys. The SET correctors are now an option that can be applied in real-time on the NavCom GPS receivers delivered to NAVOCEANO.

*Note: h_Nav and D are corrected to vertical distances on the boat common reference system, installation offsets, and lever arm corrections.

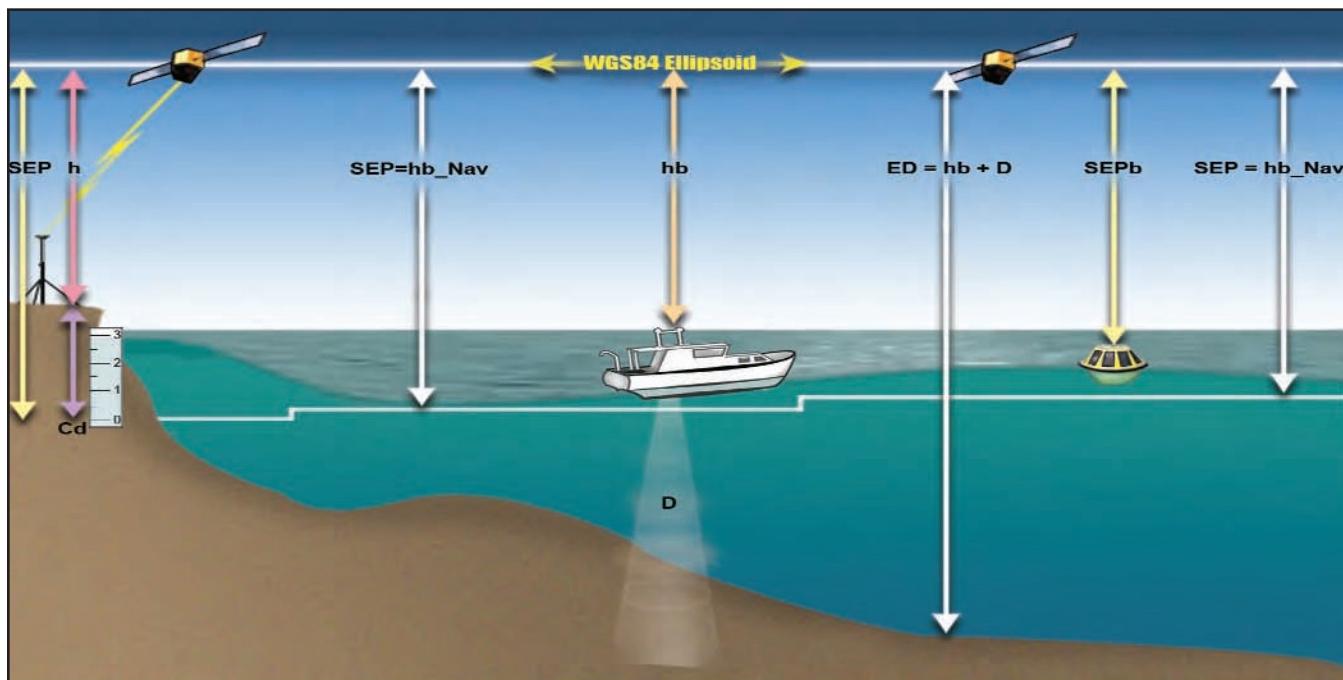


Figure 3. RTK or RTG positioning basics with option of land-based infrastructure or without, depending on survey accuracy requirements or whether access to land is granted or feasible.

RTG POSITIONING ACCURACY TESTS

The equipment configuration used in the Ship Island, Mississippi Gulf Coast, 24 February 2005, RTG positioning tests were:

- ▶ ODOM 200KHz EchoSounder
- ▶ TSS DMS20 Motion Sensor (no heading). GPS was used as Heading.
- ▶ Two NavCom SF-2050M—with the SET corrector enabled—connected via an RF splitter to the same GPS antenna.
- ▶ One operation in RTK mode (with RTK Extend Enable).

▶ The other operating in RTG+SET only. In the RTK mode, the integers were fixed. Simultaneous navigation solution was logged into HYPACK.

- ▶ Data was read, cleaned, and merged in CARIS HIPS into a depth from the ellipsoid solution.

Figure 5 shows the results of the RTG Positioning Accuracy Test performance around Ship Island.

ELLIPSOID DATUM SEPARATION ACCURACY TEST UTILIZING THE RTG/RTK BASED POSITIONING AND TELEMETRY BUOY (P&TB)

Table 1 displays the results of the comparison of the land-based Canadian Hydrographic Service (CHS) Tide Gauge in Particia Bay versus the results generated by RTG/RTK-based P&TB. The computed datums for this test are Mean Higher Water (MHHW), Mean Lower Low Water

(MLLW), Mean Tide Range, and MSL for the observation period.

THE FUTURE IS NOW—CONCEPT OF OPERATIONS FOR MILITARY SURVEYS

With the recently acquired NavCom GPS receivers (and their less than 20 centimeter (2 sigma) horizontal accuracies), meeting IHO Special Order horizontal accuracy standards worldwide is almost a trivial problem. Meeting the vertical accuracy standards, however, without the use of tide gauges ashore, requires a coherent strategy. Similar to Kinematic GPS (KGPS) surveys, soundings in RTG surveys are related in three-dimensions (horizontally and vertically) to the WGS-84 Ellipsoid.

The survey platform can be seen as an altimeter or simply as a sensor platform that accurately position depths on the GPS absolute reference frame. Charted depths become a derived product by subtracting the WGS-84 three-dimensional depth positions to a datum to ellipsoid separation (SEP) distance.

Continued Next Page...

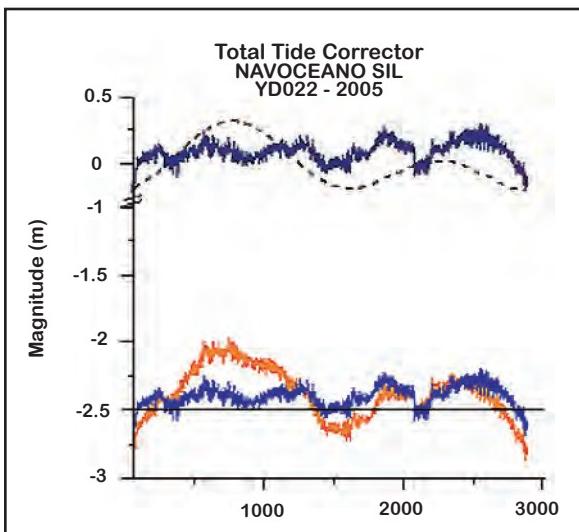


Figure 4. NavCom RTG solutions corrected with IERS Note 21 SET model. A vertical accuracy of approximately 24 centimeters at 95% confidence level over a 24-hour observation period is shown.

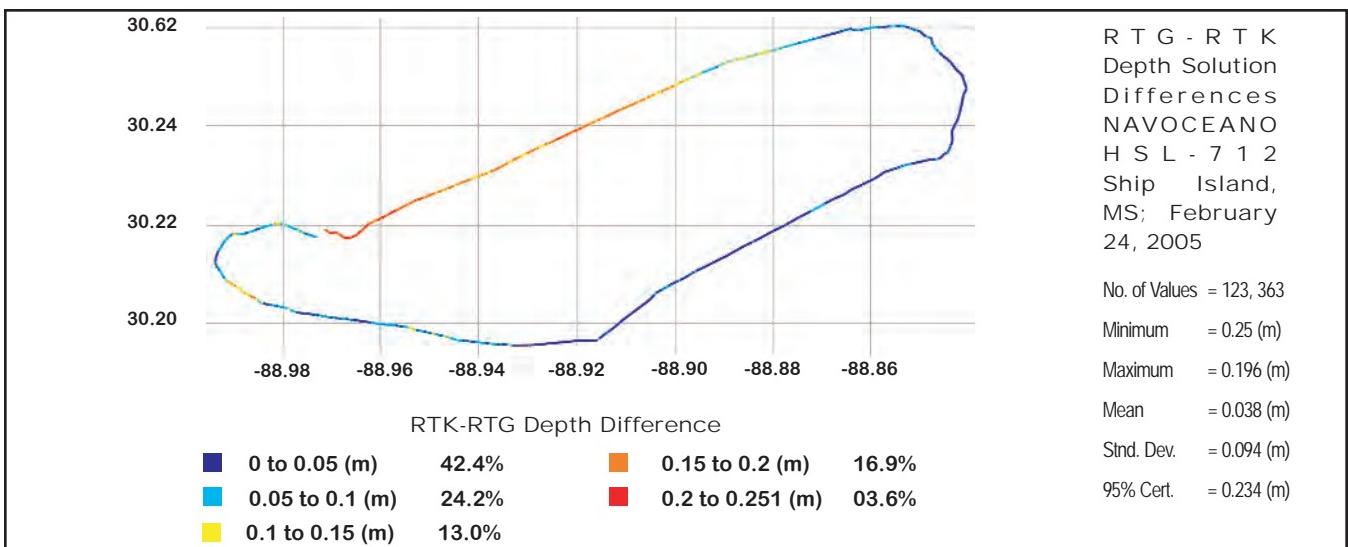
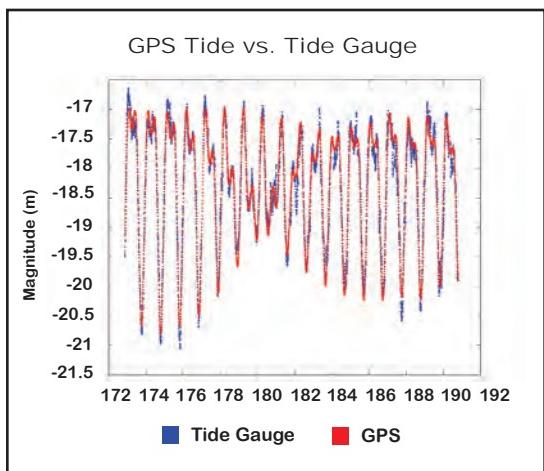


Figure 5. RTG-RTK depth solution differences for Ship Island, MS.

Figure 7 depicts a seamless bathymetry and topographic surface (ETOPO5) plotted on a computer simply by the conversion of latitude (ϕ), longitude (λ), and ellipsoidal height (h) geodetic coordinates to the WGS-84 ECEF Cartesian (XYZ) coordinate system.

Figure 8 represents a seamless equipotential surface, the EGM96 geoid and the WGS-84 ellipsoid, plotted on the ECEF Cartesian coordinate system of WGS-84. Note that for visualization purposes, the geoid and the spheroid is exaggerated 1000 times.

Charted depths can be derived and continually updated utilizing in-situ separation values and modernized Geoid and hydrodynamic models. At the moment, the following three techniques can be employed in combination:



► Traditional tide gauges accurately tied to the geocentric reference frame. This would be a preferred method when access to land is feasible and the hydrographic survey is near the tide gauge. Adding the ellipsoidal height to the tide benchmark will produce a chart-datum SEP value than is valid if the amplitude, the range, and slope of the MSL do not change much within the survey area. When the survey is far from the tide gauge, a SEP must be a complement-built model to correct for the absolute depth measurements. SEP models can be built using a combination of other sensors and information. (See Figure 5.)

► RTG/RTK GPS-equipped buoys. By observing the water elevations and tide ranges the local MSL (with respect to the WGS-84 Ellipsoid) can be determined. With at least 15 days of data, the four major constituents can be derived from the tidal signal. The advantages of buoys are that they are easy to install, data can be easily retrieved, and tide levels can be directly related to WGS-84. The disadvantage of this technique is, like all small buoys, they are susceptible to tampering and severe

weather. Initially, the buoys are similar in cost to bottom gauges, although prices are expected to decrease. (See Figure 6.)

► Bottom-mounted tide gauges. By observing the water elevations and tide ranges the local mean sea level (with respect to the gauge) can be determined. With at least 15 days of data, the four major constituents can be derived from the tidal signal. Soundings taken over the top of the gauge with a launch will provide a vertical reference to the WGS-84 Ellipsoid. The disadvantages of this technique are the difficulty in installation and retrieval and the delay in obtaining needed data until the end of the observation period. Bottom gauges are also expensive and susceptible to bottom trawls by fishing vessels. But bottom gauges, if corrected for barometric pressure and properly installed, provide very reliable data.

► Numerical tide models. Numerical models to determine MSL, Chart Datum, and the WGS-84 Ellipsoid relationships can also be used. Tide models are built using the best available bathymetry, coastline, known tide stations, possibly weather information, and boundary conditions. With these data the models, using the equations of motion, determine tide heights at

Figure 6. CHS tide gauge data (red) and NAVOCEANO P&TB data (blue).

	CHS TIDE GAUGE (M)	P&TB (M)	DIFFERENCE (M)
MHHW (M)	-17.64	-17.57	-0.07
MLLW (M)	-20.55	-20.59	0.04
MEAN TIDE RANGE (M)	2.91	3.02	-0.11
MSL (M)	-18.78	-18.79	0.01

Table 1. NAVOCEANO Tide Analysis System (NAVOTAS) tide record analyses for Tide Gauge and P&TB data.

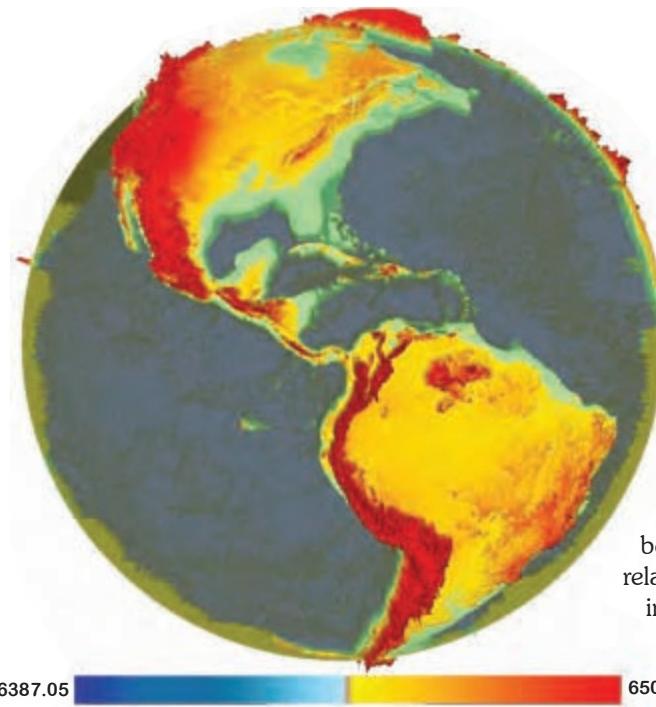


Figure 7. ETOPO5 on the WGS84 ECEF Cartesian coordinates system. Color code reference from EGM96 GEOID.

Datum to WGS-84 Ellipsoid relationship (SEP).

Ideally, the modeling technique will be used to define the general characteristics of the Chart Datum/WGS-84 model and determine any steep slopes in the area. Two to three GPS buoys will be used to refine the modeled relationships throughout the immediate survey area or a continued

re-deployment of the GPS buoy from Survey Operation (SURVOP) to SURVOP. A typical SURVOP lasts approximately 21 days.

CONCLUSION

NAVOCEANO can no longer rely on traditional methods for conducting hydrographic surveys. The terrorism threat and nations unwilling to provide access to shore sites for hydrographic survey stations has severely decreased traditional NAVOCEANO shore-dependent hydrographic survey operations.

Furthermore, these threats have shifted the NAVOCEANO focus to providing the warfighter with a high-resolution, near-real-time depiction of the battlespace environment. The result is the NAVOCEANO adoption of the NavCom Technology, Inc., GPS receivers with RTK and RTG capabilities and associated StarFire™ WADGPS.

While the primary motivation for NAVOCEANO to

implement this highly accurate mapping technique was to solve a security and logistic problem, the consequences are highly significant to the scientific community within and outside NAVOCEANO.

These highly accurate seabed and sea surfaces are referenced to the same absolute geocentric system as other highly accurate earth measuring systems such as the GRACE and CHAMP gravity satellites, satellite altimeters, and airborne LIDARs among others.

This opens numerous opportunities and applications to oceanographers, geodesists, and solid earth scientists to study global seabed and ocean dynamics on a finer scale than has been previously possible.

This will, in turn, enable a better understanding of the processes that drive the Earth's dynamic system (e.g., solid Earth, ocean, and atmosphere), thus leading to better analysis and predictions of climate change and natural hazards (e.g., earthquakes).⁴

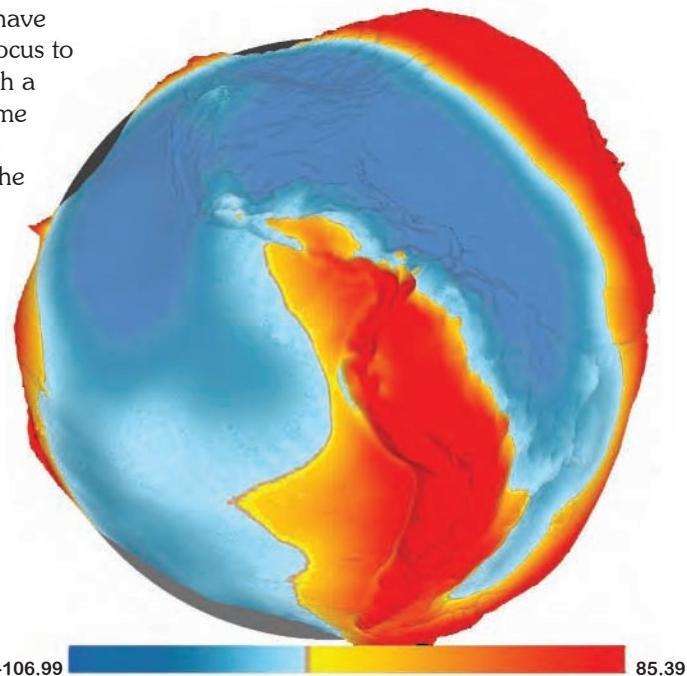


Figure 8. Seamless EGM96 GEOID and WGS-84 spheroidal surface plotted on cartesian coordinates using WGS-84 ECEF reference system. Red-orange tones represent geoidal ondulations above the reference ellipsoid. Blue tones represent geoidal ondulation below the reference ellipsoid.⁵

References

1. Young, R. "Law of the Sea," Multibeam and Visualization Workshop, 26-29 July 2004, Gulfport, MS.
2. International Hydrographic Organization Standards for Hydrographic Surveys, Special Publication 44, 4th Edition, International Hydrographic Bureau, Monaco, April 1998.
3. Hatch, R., T. Sharpe, and P. Galyean, "StarFire: A Global High Accuracy Differential GPS System," NavCom Technology, Inc., Paper 1.6, 30 October 2002.
4. Tapley, B. D., S. Bettadpur, M. M. Watkins, and C. Reigber, "The Gravity Recovery and Climate Experiment: Mission Overview and Early Results," Geophysical Research Letters, 31, L09607, doi:10.1029/2004GL019920, 2004.

Background Materials

Arroyo-Suarez, E., G. Glang, J. Riley, and D. Mabey, "Evaluating a Global Differential GPS System for Hydrographic Surveying," Proceedings of the Institute of Navigation Global Navigation Satellite System (GNSS) Conference 2005, 21-24 September 2005, Long Beach, California.

Arroyo-Suarez, E., and V. Hsiao, 2004: "Naval Oceanographic Office Positioning System and Intentions to Operationally Geolocate Bathymetric Measurements to a Seamless Reference Frame" Proceedings of the Institute of Navigation GNSS 2004. 21-24 September 2004, Long Beach, California.

Center for Operational Oceanographic Products and Services, Computational Techniques for Tidal Datums Handbook: NOAA Special Publication NOS CO-OPS 2, Silver Spring, MD, U.S. Department of Commerce, National Oceanic and Atmospheric Administration, National Ocean Service, September 2003.

Van Norden, M., S. Harrison, and S. Kosbab, "Electronic Chart Display Overlay Development at the Naval Oceanographic Office for use in Tactical Applications," 2nd International Electronic Chart Display and Information System (ECDIS) Conference and Exhibition, 7-9 October 2003, Singapore.

Van Norden, M., and E. Arroyo-Suarez, "Concept for Military Surveys to IHO Standards Without Shore Stations Using the Real-Time GIPSY (RTG) Global Positioning System," Proceedings of the U.S. Hydrographic Conference, 29-31 March 2005.

Disclaimer

The mention of any commercial product or company name in this paper does not constitute an endorsement by the Naval Oceanographic Office or the U. S. Navy Fleet Survey Team.

Practical Parallel Graphics...continued

applications. Known as Secure Shell Wrappers, this was presented at the 2005 User's Group Conference (UGC). Secure Shell Wrappers use a standard SSH (or Kerberized Remote Shell) connection to start the server and also act as the data channel. This is less work for the user, and better conforms to the secure access guidelines for the High Performance Computing Modernization Program (HPCMP)—though it does require modifications to the visualization program. Secure Shell Wrappers can be used in conjunction with MPI, Chromium, or VGP to do parallelized, remote visualization.

Beyond Secure Shell Wrappers and Paraview, several other options are under investigation. One option is Paraview Enterprise Edition, a commercial version of Paraview that will offer Web-based access.

Another commercial product under investigation is RealityServer, which provides Web-based access for custom applications.

Finally, there is Chromium with Virtual Network Computing (VNC), which should make it possible to access the server with only a VNC client and little or no modification to the visualization program.

MORE INFORMATION

It can be difficult to decide upon the best remote and parallel visualization approach given a specific application. At the NAVOCEANO MSRC VADIC, we recommend examining Paraview first (at <http://www.paraview.com>). For more information and support, a Paraview course that specifically focuses on remote and parallel visualization is planned at UGC 2006. In addition, visualization staff from

NAVOCEANO and other MSRCs also plan to present topics in remote and parallel visualization at UGC 2006.

For more difficult decisions involving complex data, deciding between Paraview and custom applications, or for other problems or questions, users may contact the user support or visualization support staff at their respective MSRC or the NAVOCEANO MSRC VADIC visualization support staff at viz@navo.hpc.mil.

After deciding upon a visualization approach, users may need to configure their accounts and tailor their environment. To learn about NAVO MSRC's scientific visualization resources, configuration, center-specific Paraview instructions, and other helpful information, see the NAVOCEANO MSRC Web site at <http://www.navo.hpc.mil/>.

Five Dogs...continued

up the pieces. Debit card no longer being honored; Capital One won't increase my credit limit. Low on cash, and there's gas shortage to add to the drama. Keesler Credit Union down due to Katrina—financial status: bleak.

- ▶ Wed., 31 Aug.: Spent day searching for and buying small generator, window air conditioning unit, gas cans, camping supplies, food, etc.
- ▶ Thur., 1 Sep.: Liberated the frisky dogs from the kennel and drove to Jackson, MS, for Phase 1 of the trip home. Power out in much of Jackson and unable to buy gas.
- ▶ Tue., 6 Sep.: Gas availability improved so continue trek to Diamondhead, MS, home of the two beagles. Set up generator and window air conditioning unit in the beagle's house only to find out from a neighbor that power was restored to the house, but a Power Company representative had turned the circuit breaker off at the meter. Flipped the breaker "on."
- ▶ Wed., 7 Sep.: Reported briefly to the NAVO MSRC to get on the Federal Emergency Management Agency (FEMA) list for mobile homes; unpacked cars and set up base camp in

beagle's guest bedroom. Gas and food shortages. Discover that DoD "Meals Ready to Eat" (MREs) are filling but not especially tasty.

- ▶ Thur., 8 Sept. - Sat., 10 Sept.: Tunneled through house to rescue my wife's completed needlepoint works and china. (Everything is a loss.) Water level rose higher than 10 feet in house—ceiling collapsed, several inches of sewage smelling water and mud fills the house.
- ▶ Sun., 11 Sep.: Exhausted—took the day off.
- ▶ Mon., 12 Sep.: Returned to work; bought a chocolate donut on the way to work—life is good again!

"Guttin" my hurricane-damaged house made the second phase of the TI-06 process a challenge as there were times when I had to stop debris removal and house-cleaning work altogether (including weekends) to

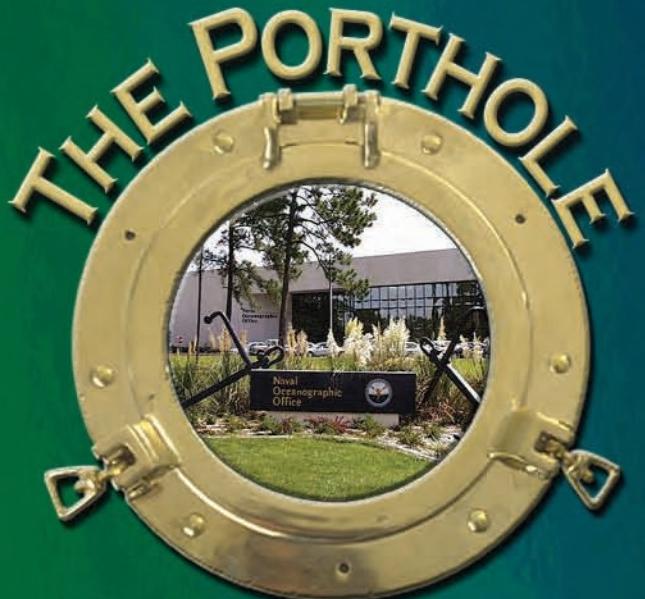


keep on schedule. It was definitely worth it, though, as I was able to present the Usability Phase II results to the CAT on schedule.

I do need to note, though, that throughout the ordeal, the "U" Team members (see below) made my task easier by providing timely and thorough evaluations of the proposed systems along with much appreciated encouragement. I am indebted to them for their support; in particular to Tom Crimmins for pinch-hitting for me during the Phase I Usability presentation to the CAT.

- ▶ Tom Crimmins and Mike Knowles, ARL
- ▶ Ralph McEldowney and John Gebhardt, Aeronautical Systems Center (ASC)
- ▶ Bobby Hunter and Paula Lindsey, U.S. Army Engineer Research and Development Center (ERDC)
- ▶ Lee Higbie, Arctic Region Supercomputing Center (ARSC)
- ▶ Tim Fahey, Maui High Performance Computing Center (MHPCC)
- ▶ Steve Schraml, User Advocacy Group (UAG)
- ▶ Ed Farrar, Naval Oceanographic Office (NAVOCEANO)

The Dining Room: Pre-Katrina (top), post-Katrina (left), and today (right). The sideboard in the "pre" picture floated across the room to the position seen in the "post" picture.



visitors and events at NAVO MSRC



NAVO MSRC Director Steve Adamec explains the MSRC supercomputing capabilities to Northrop Grumman VIPs.



Northrop Grumman VIPs receive a demonstration of MSRC scientific visualization capabilities.



NASA visitors (L-3 Communications) receive a demonstration of MSRC scientific visualization capabilities.



Semeion A. Sertsu and LCDR Lora Turner, Executive Officer of Naval Ice Center, tour the NAVO MSRC with the MSRC Deputy Director Bobby Knesel.



RADM Fred Byus,
Oceanographer/Navigator
of the Navy,
visits the MSRC
for a familiarization
tour of the Navy's
oceanographic activities.



Naval Meteorology and Oceanography (METOC) students visit the MSRC.



NAVOCEANO recipients of a Team Award for members aboard the USNS Henson, flanked by CAPT Andrew Brown III and CAPT John Cousins.



BMC Cecil Watt takes a quick lunch break in the shade with an athlete he supported during the Special Olympic games.



Brenda Villalobos of NAVOCEANO takes a photo of a Special Olympian and volunteers that supported her in the games.



Boy Scouts of America National Jamboree 2005

Every four years the Boy Scouts of America (BSA) hold a national Jamboree. As in the past, Naval Oceanographic Office (NAVOCEANO) personnel joined with Navy and United States Coast Guard (USCG) personnel to participate. NAVOCEANO's Don P. Ecuyer provides a brief narrative of some of the NAVOCEANO participation in the National Jamboree 2005 (JAMBO 2005).

As in 2001, I was once again asked to participate in the BSA JAMBO 2005. Along with Steve Sramek (NAVOCEANO), I was a member of the Oceanography Merit Badge staff with a member of the USCG as the head. As a Scouter (an adult involved in scouting), I am qualified to teach several merit badges including oceanography.

When we arrived at the Merit Badge Midway of JAMBO 2005, we learned the USCG lead had to cancel because of family illness and I was in charge. All we had was an empty booth, the few materials we had gathered from NAVOCEANO, and 96-degree heat. With the help of the Navy contingent from the Naval Reserve Recruiting Command in Millington, Tennessee, we were able to put together a presentation that was equal to any of the merit badge booths on the Merit Badge Midway.

Lord Baden-Powell, the founder of the Boy Scouts, once said that boys learn when they are having fun. We focused on this as we planned our merit badge presentation: show the boys interesting items related to oceanography, and their attention will be focused on the merit badge. We did this using the NAVOCEANO-supplied

materials and by focusing on interesting subjects taken from the merit badge book.

At each jamboree, I am always surprised by the interest in oceanography. When we opened the booth on the first day, there were three groups of boys waiting to sign up. They were from Kansas, Utah, and South Dakota. The scouts from Utah said that oceanography was the merit badge they were most looking forward to because it was unavailable in their area.

During JAMBO 2005, nearly 400 boys took the Oceanography Merit Badge course. The course began with a presentation, which adheres to the requirements needed to complete the merit badge. By attending the three-hour lecture, the scouts were able to complete seven of the nine requirements. The last two requirements were finished by the scouts during their free time: a 500-word paper on oceanography and a small project. More than 90 percent of the 400 scouts who took the merit badge completed the requirements. Most of the time, when a scout signs up for a merit badge with a counselor from his local BSA council, he usually is taught by someone who has learned the subject matter as a hobby. At the JAMBO, the scouts are exposed to counselors who have expertise in their fields usually from professional experience. The four main volunteer instructors for oceanography were myself and Steve Sramek, a person from USCG Environmental Protection, and an oceanography professor from Texas A&M University.



In a JAMBO 2005 highlight, President Bush addresses JAMBO participants.



Don P. Ecuyer (NAVOCEANO) works with a few of the nearly 400 scouts who worked toward an Oceanography Merit Badge.

Navigator Tools and Tips

Hints for Choosing “Consumable CPUs” and “Consumable Memory” Resource Requests on KRAKEN and ROMULUS

Sheila Carbonette, NAVO MSRC User Support

Choosing the correct values for the keyword “#@ resources” can be one of the most complicated tasks when creating a LoadLeveler batch script.

If the wrong values are chosen, the AIX Workload Manager (WLM) will terminate the batch job and send e-mail to the user with a brief description of why the job was killed.

The WLM was implemented on the IBM P4+ systems, ROMULUS and KRAKEN, to help control resource utilization during periods of peak system demand. It monitors the system resources to prevent jobs from interfering with each other when there are conflicting resource requirements. This is why it is very important to choose the correct values for the “Consumable CPUs” and “Consumable Memory” LoadLeveler keywords.

Now for hints in choosing the resources: choosing the value for “Consumable CPUs(#)” is pretty straightforward. For strictly non-parallel serial jobs and Message Passing Interface (MPI) parallel jobs, the “Consumable CPUs(#)” is always set to 1.

For OpenMP executables, the “Consumable CPUs(#)” is set to the maximum number of threads (processors) your program will use. This value must match the value for the “OMP_NUM_THREADS” environment variable.

Calculating the value for “Consumable Memory” is a little more complicated. If you know how much memory your job needs for each node, the value can be calculated by dividing the total amount of memory in the Memory Buffer (MB), by the number of tasks requested for the node.

If you are not sure how much memory your job needs, you can start with the maximum available for each node. Refer to Tables 1 and 2 for the maximum amount of consumable memory available for each node based on LoadLeveler class on KRAKEN and ROMULUS.

There are also other LoadLeveler commands and hardware performance utilities available to estimate the amount of memory your job needs. The LoadLeveler command, “lacct” with the “-jobid” option can be used to display the highwater memory mark of a running job. The output of this command is also written to *stderr* when the job finishes. If the highwater memory mark of real-memory is lower than you requested, you can lower the value specified for the “Consumable Memory.” If the highwater memory mark of real-memory is higher than you requested, you may need to ask for more processors.

Continued Next Page...

KRAKEN			
LOADLEVELER CLASS	MAX. NUMBER NODES	CONSUMABLE CPUs	MAX. CONSUMABLE MEMORY
background	24	24	1292 MB
bigmen	04	32	58624 MB
block	16	128	12592 MB
challenge	128	1024	12592 MB
debug	32	256	12592 MB
high	128	1024	12592 MB
share	01	01	2048 MB
standard	64	512	12592 MB
transfer	01	01	2048 MB
urgent	128	1024	12592 MB

Table 1. Maximum amount of consumable memory available for each node based on LoadLeveler class on KRAKEN.

The hpmcount utility available in the Hardware Performance Monitor (HPM) Toolkit can be used to monitor the resources and performance of your application. The HPM ToolKit can be found in the /opt/HPC_ToolKit directory on each of the IBM systems. The utility also displays the highwater memory mark of the job.

For serial jobs the following can be used to start the utility:

```
hpmcount -o outfile executable_name
```

where “-o outfile” is the name of the output file that is generated.

For MPI jobs use the following:

```
poe hpmcount -o outfile executable_name
```

In this case, the "-o outfile" will generate an output file for each process with the following naming convention: "outfile.<pid>".

If you are still unsure of your job's memory requirements, please contact NAVO MSRC User Support for assistance.

Below are “#@ resources” examples.

A serial job running in the standard queue:

```
#@ job_type = serial
#@ node = 1
#@ tasks_per_node = 1
#@ resources = Consumable CPUs(1)
Consumable Memory(12592 mb)
```

The resources specified above will give this serial job access to one processor and a total of 12592 megabit (Mb) of memory.

A serial job running in the share queue:

```
#@ job_type = serial
#@ node = 1
#@ tasks_per_node = 1
#@ resources = Consumable CPUs(1)
Consumable Memory(1574 mb)
```

The resources specified above will give this serial job access to one processor and a total of 1574 Mb of memory.

A serial job running in the transfer queue:

```
#@ job_type = serial
#@ node = 1
#@ tasks_per_node = 1
#@ resources = Consumable CPUs(1)
Consumable Memory(512 mb)
```

The resources specified above will give this serial job access to one processor and a total of 512 Mb of memory.

An OpenMP job running in the standard queue:

```
#@ job_type = serial
#@ node = 1
#@ tasks_per_node = 1
#@ resources = Consumable CPUs(8)
Consumable Memory(12592 mb)
setenv OMP_NUM_THREADS 8
```

The resources specified above will give this OpenMP job access to eight threads and a total of 12592 Mb of memory.

An MPI job running in the standard queue:

```
#@ job_type = parallel
#@ node = 1
#@ tasks_per_node = 8
#@ resources = Consumable CPUs(1)
Consumable Memory(1574 mb)
```

The resources specified above will give this MPI job access to eight processors per node and a total of 12592 Mb of memory.

ROMULUS			
LOADLEVELER CLASS	MAX. NUMBER NODES	CONSUMABLE CPUs	MAX. CONSUMABLE MEMORY
background	24	192	12592 MB
bigmen	04	32	58624 MB
block	46	368	12592 MB
challenge	24	192	12592 MB
share	02	1	2048 MB
standard	24	192	12592 MB
transfer	01	1	2048 MB
urgent	24	192	12592 MB

Table 2. Maximum amount of consumable memory available for each node based on LoadLeveler class on ROMULUS.

UGC2006 - Users Group Conference
Denver, June 26 - 30

[http://www.hpcmo.hpc.mil/
Htdocs/UGC/index.html](http://www.hpcmo.hpc.mil/Htdocs/UGC/index.html)

SERA 2006 - 4th ACIS Conference on Software Engineering Research, Management & Applications
Seattle, June 26 - 30

[http://www.hpcmo.hpc.mil/
Htdocs/UGC/index.html](http://www.hpcmo.hpc.mil/Htdocs/UGC/index.html)

SIP 2006 - 8th IASTED International Conference on Signal and Image Processing
Honolulu, Aug. 14-16

<http://www.iasted.org/conferences/2006/hawaii/sip.htm>

SCC 2006 - IEEE International Conference on Services Computing
Sept. 18-22 Chicago

[http://conferences.
computer.org/scc/2006/](http://conferences.computer.org/scc/2006/)

Cluster 2006 - IEEE International Conference on Cluster Computing
Barcelona, Sept. 25-28

<http://www.cluster2006.org/>

Gridnets 2006 - 3rd International Workshop on Networks for Grid Applications
San Jose, CA, Oct. 1-2

<http://gridnets.org/2006/>

ICCD 2006 - 24th International Conference of Computer Design
San Jose, CA, Oct. 1-4

[http://www.iccd-
conference.org/
generalinform
ation.htm](http://www.iccd-conference.org/generalinformation.htm)

VIS 2006 - 2006 IEEE Visualization Conference
Baltimore, Oct. 29 - Nov. 3

[http://vis.computer.org/
vis2006/](http://vis.computer.org/vi s2006/)

SC06 - International Conference for High Performance Computing, Networking & Storage
Tampa, Nov. 11-17

<http://sc06.supercomp.org/>

ISSRE 2006 - 17th IEEE International Symposium on Software Reliability Engineering
Raleigh, NC, Nov. 6-10

[http://www.csc2.ncsu.edu/
conferences/issre/](http://www.csc2.ncsu.edu/conferences/issre/)

Coming Events
Coming Events

$$y = R \times \cos(\text{lat}) \times \sin(\text{lon})$$

Naval Oceanographic Office * MAJOR SHARED RESOURCE CENTER
1002 Balch Boulevard . Stennis Space Center, Mississippi . 39522